Council of European National Top-Level Domain Registries

Report on IETF103

Bangkok
3-10 November 2018

CENTR vzw/asbl · Belliardstraat 20 (6th floor) · 1040 Brussels, Belgium
Phone: +32 2 627 5550 · Fax: +32 2 627 5559 · secretariat@centr.org · www.centr.org

Contents

Highlights

Human Rights Reviews not welcome in some spaces at the IETF

The Human Rights Protocol Considerations Research Group (HRPC RG) has started to carry out more regular reviews on emerging protocols from a human rights' perspective (looking for privacy issues, but also checking things like accessibility and internationalisation or market concentration effects). During the IETF week in Bangkok, two reviews in particular led to fierce discussions which left the NGOs performing the reviews baffled. Former EU Member of Parliament, Amalia Andersdottir, who is now working on Article19, spoke about the "mixed signals" that were hard to explain. Some working groups (for example Suit or IPWave) welcomed the reviews, but IETF Chair Alissa Cooper and outgoing Security Area Director and IESG member Eric Rescorla were very disparaging about the report performed on QUIC.

HRPC RG report on QUIC felt to be unhelpful by IETF Chair

The Human Rights review of the new QUIC protocol received its worst "review" during the session of the Human Rights Protocol Considerations Research Group (HRPC RG) at IETF103.

The review had added nothing to discussions in the QUIC working group, said Transport Area Director, Mirja Kühlewind (ETH Zurich). The review was full of technical errors and was too far removed from the standardisation work, said Security Area Director Eric Rescorla. IETF Chair Alissa Cooper added that the review essentially lacked direct interventions by concerned experts in the WG and pointed instead to the "great work" the IETF itself had done to enhance privacy in recent years.

Cooper also observed that discussions on "Privacy Reviews" for draft workshop documents which had once been proposed had gone nowhere. The Privacy Reviews, an idea based on pre-Snowden work on "Privacy Considerations for Internet Protocols" (RFC 6973), were intended to be undertaken by members of the IETF Security Area. However, as the Security Area has been much less active recently, partly due to workload of the Chairs (Eric Rescorla is an active draft author and has recently become Mozilla's CTO), formal privacy reviews have not been carried out in a long time.

Niels Ten Oever, co-founder of the HRPC RG and co-author of the QUIC (as well as other) reviews welcomed the feedback, announcing that it would help inform the work on a guideline document about how human rights reviews should be performed in the future. The document lists five methods (also used for the QUIC review):

- Analysing the drafts based on the guidelines on the human rights considerations model
- Analysing the drafts based on their potential impact
- Expert interviews
- Interviews with impacted communities
- Tracing impacts

Ten Oever suggested that the interviews had certainly resulted in mirroring working group discussions and announced that this could be reconsidered. Talking to the experts on the other hand would allow the HPRC RG to stay as on top of the technical details of a draft as possible.

The basic question for the HRPC RG now is how to go forwards. Given the rather harsh comments, it might become more difficult to find reviewers. At the same time the engineers working on the draft standards themselves are often too busy to get the drafts done. Can the IETF manage without the external review on privacy and other impacts of technology which standard practice? Cooper seems to hint that the organisation can. Nevertheless Snowden's revelations and the reaction of the IETF to this in its intensifying privacy work suggests otherwise.

Hot debate on human rights review of RDAP verification extension in the RegEXT WG

The broadly discussed QUIC documents were not the only drafts that human rights reviewers took on. This time the Registry Extensions WG was provided with an assessment with regard to the RDAP verification extension. The HPRC report listed a number of points for the WG to consider, the first three being:

With regard to privacy: VSP (verification service providers who act as a third party) are obliged to ("MUST") collect and store personally identifiable data (domain name, registrant contact).

With regard to content control/censorship: VSP will check "whether the domain name is not prohibited or whether the registrant is a valid individual organisation, or business in the locality".

IPR – since the IETF favours open standards, parts of the specification are covered by a patent applied for by Verisign. According to the review "this includes a description of the grace period within which the requirement set of verification codes may be sent before the object becomes non-compliant", and "a clear depiction of the flow of the request detailed in Fig. 1 of the [PATENT]."

In conclusion, the reviewers decided that the draft standard had issues and recommended that the WG should at least be transparent about the risk in a human rights consideration section. Such a section would be added, along with standard security and optional privacy to the end of the specification text. They also underlined that the WG could consider not making it a full standard from the start, as so far only one company (Verisign) is implementing it. In that way the extension essentially does not fulfil IETF process, as it usually calls for several independent implementations before being considered as a "standard" qualification.

After a controversial previous discussion on the mailing list, WG Chair Jim Galvin set some time aside during the Bangkok meeting. But the two sides did manage to make steps towards potential consensus over the issues. Document editor and author James Gould acknowledged one issue, namely regarding the PII data retention obligation for VSPs and made a change in the declaration of this subject to local legislation. VSPs need to store their validation decision for a domain registrant. Additional data (name, address of registrant) storage is subject to local privacy legislation.

Apart from that, Gould said he did not feel competent enough to add a "human rights consideration section" which, according to the proposal would highlight the potential risks for privacy, discrimination and accessibility. His Verisign colleague Scott Hollenbeck questioned the idea of "standard" human rights considerations in draft RFCs. There was no consensus policy in the IETF on this and the RegEXT WG should not be made the test case for that, he said.

RegEXT Chair Galvin underlined several times that the considerations could only be treated as individual contributions, the same way as any other technical contributions to tech issues. Galvin even went as far as to say that the WG should focus on technology only, as policy issues were outside of the process. That statement was sharply challenged by Ten Oever, who said that calling technology and standardisation neutral rather than political was a naive conception.

Reverse Search: the next candidate for a privacy/ human rights review?

Perhaps a better view on how standards are policy in part was delivered by the WG itself in its discussion on the RDAP reverse search extension. Loffredo (.it Registry) presented the extension, which would allow a reverse search over the RDAP database starting from various data points, as a candidate for a new milestone. The reverse search extension is mainly dedicated to registrars, allowing them to search for their own domains, and, it should be provided under a strict control based on user access levels. After clearing most of its milestones (see below), the WG has re-chartered and is taking on new extension documents for the next iteration.

When presenting the proposal Loffredo himself pointed to ICANN policies, and more specifically to two documents produced by ICANN: the next-gen RDS (2014) and the Registry Agreement specification (2017). Chair Galvin pointed out that there are lots of discussions going on regarding the RDAP rollout, and the WG work cannot be motivated solely by ICANN policies. WG member Wilhelm (Verisign) warned that one of the documents cited (that is, the next-gen RDS) had been closed down and that the draft was very early in RDAP implementations, given that ICANN's community was still working on post GDPR registration policies. Therefore, these things create an implementation burden for contracted parties. While Alvarez (ICANN) requested reverse search capabilities during the session, an unchecked reverse search capability could potentially have an impact on privacy aspects. Finally, Loffredo replied that the draft was not influenced solely by ICANN policies and the draft's authors were thinking about taking a controlled approach to reverse search.

The main takeaway with regard to the relation of standards and policy is that on occasion, it is both ccTLDs and gTLDs policies that are cited as providing the legitimation for the standards work. The RegEXT is perhaps one of the groups that best illustrates how technical standards work treads a fine line between standards and policies.

The decision about taking on the reverse search as well as the other proposals currently asking for adoption under the new charter (see below) will be taken either on the mailing list or during the next RegEXT.

It is unclear whether or not the RegEXT WG will come back to the recommendations of the HRPC RG. Ulrich Wisser from ISS argued during the WG discussion that making the potential privacy issues transparent in short section "was little enough". It also is a safe bet that the HRPC members present at the WG will come back with more reviews (possibly on reverse search). The RegEXT WG, which is a relatively small WG, will certainly call for more review of its documents, which have an impact on millions of registrants globally.

Rough Consensus? The IETF Plenary hears complaints about the aggressive tone in WG discussions

It took just a brief statement from one of the people responsible for mentoring newcomers to the IETF to open the flood gates to an extended me2-like discussion over the sometimes-aggressive behaviour in IETF working groups. He had been made aware of many shocking stories, Wes Hardacker said, as he read a carefully crafted statement in the plenary meeting. While he remembered being fascinated by the immense passion at meetings, he had understood later that it came at the expense of others - and he remembered how nervous he had been the first time he spoke at an IETF WG. Hardaker called on IETF participants to compose their phrasing with care, even when speaking to more seasoned IETF engineers, and regardless of how thick they might think those participating in the debate's skin might be.

It was only after Adam Roach, WebRTC Area Director and observer of the discussions at RegEXT, chimed in and said that there was now more sensitivity for the issue than at earlier times that mike lines started growing and some people warned that the problem, instead of being addressed, had got worse.

Resolverless DNS and other next steps from DoH

Given the hot discussions that DNS over HTTPs and Mozilla's ongoing test have elicited, it was rather quiet around that topic at the IETF in Bangkok. Still the DNS has become a contested area, it seems.

Patrick McManus said to this reporter that he felt that Mozilla's test operation had been mis-characterized and that next steps were still to be discussed. A dozen DoH servers (including Cloudflare, Quad9, Google, PowerDNS) are currently publicly available according to the GitHub DOH side project. The DNS Privacy project keeps track and notes that Chrome was "working on exposing DoH via a user configuration option with a drop down list and user defined option".

The DoH WG with the publication of RFC 8484 "DNS Queries over HTTPS (DoH)" went dormant, waiting to decide if it should re-charter. According to one of the Chairs, potential follow-up work could be handled by DNSOP, DPRIVE, or httpbis. In some instances, the blending of DNS and HTTP expertise might be required though. One request on the mailing list came from Bert Hubert (Power DNS), who said that in its current form DoH is ineffective, as users need "around 22 packets per DNS query/response". Hubert noted that TLSv1.3 might improve this and that "a 'slightly suboptimal' network absolutely kills browsing performance in Firefox Nightly using DoH" (0.5% packet loss turns into a 5% failure rate per DoH query). Hubert, one of the critics of DoH called for considerations for a draft on a DoH3 version.

Paul Hoffman (ICANN) presented one proposal on how the choice for a DoH resolver could be organised for a DoH user during the DNSOP working group session. Hoffman called it early days for the proposal. At least one developer from the "web camp" told this reporter that other concepts, like the one proposed in the DRIU BoF at IETF102 were more likely to be pursued. DRIU looks into randomly choosing DoH servers (the bloom filter concept).

At the same time an idea to rely on servers to push additional answers (for some DNS records) could become a part of upcoming proposals.

A side-meeting on resolverless DNS

Another idea to possibly draw DNS answers into the web world is the idea of resolverless DNS, which basically considers how to answer queries using not only the DNS record that is sought after, but also additional DNS information that is cached at the local client (one asks for example.com and is offered foo.example.com – depending on the scope and the domains outside the zone which is originally asked for). The original idea came from Daniel Kahn Gilmore (ACLU) who proposed it as a way to avoid tracking DNS requests.

While no formal proposal has been made at the IETF, and the resolverless mailing list has been very quiet, the side-meeting essentially ended with three potential solutions:

1. DNSSEC-signed, any outbound link or resource from a site
2. CDN-local links or resources
3. Within TLS: any automatically loaded resource

The proposals implicitly point to concerns over the various mechanisms. Concerns include the loss of control where traffic is directed (as a typical DNS resolver is no longer in the loop) and more security concerns. If DNSSEC is not deployed and adapted to the scheme, redirecting traffic would be an issue.

Some think that different standards for accepting records might be used. Domains not covered by a certificate of the website pushing them could be made dependant on DNSSSEC validation.

To allow for some control, the idea that hostnames themselves should have the possibility to opt-in or opt-out of having their names pushed by other services was discussed.

Nevertheless, there are still concerns over replay attacks, load balancing and similar-looking domains. An attacker could for example be pushed by a server, so that fakebook.com could be served instead of facebook.com. Furthermore, load balancing would be interfered with and replay attacks could be nurtured.

The consequences for DNSSEC (the need to be deployed at browser level and possibly disincentivising OS from deploying) was discussed as well.

So far the discussions have been inconclusive. On the resolverless mailing list, nascent ideas such as "a new HTTP request header, for example 'Accept-DNS'" (Justin Henck, Jigsaw) have been put forward, but it remains to be seen if a formalised draft will be presented at the IETF. Ben Schwartz, also Google Jigsaw and DoH WG Co-Chair, said that the resolverless concept was too speculative so far to merit being adopted in the DNSOP or the DoH WGs, and that he "did not see any way to do this in the web security model".

Hyperlocal going into the next round: RFC 7706 bis

Cutting response times, adding resilience and also adding some privacy were the aims of RFC 7706. Paul Hoffman (ICANN) presented the WG with a proposal to do a bis-version for what has been called "hyperlocal rootzone" at ICANN.

The RFC 7706 by Hoffman and Kumari (Google) was adopted in 2015. The basic idea is to fetch a copy of the root zone file to avoid sending queries up to the root servers. Instead, queries can be answered locally (from a loop back server, to avoid answers being offered outside the local network).

Several root servers already offer the option for third parties to download the complete root file.

Some open source DNS software has implemented the hyperlocal concept already. Ondrej Sury (BIND) reported that BIND would (whilst cooperating with ICANN) include a local copy of the root in the next version. Unbound also uses the 7706 concept. Knot has experimented and found that it could be more effective than current solutions, according to Petr Spacek (CZ.NIC). Spacek reported that Knot currently uses a combination of NSEC 3's aggressive caching and pre-fetching).

Hoffman said that in the new bis-draft he would focus on discussing whether the root server needed to be on a local machine, or could be operated outside by a third party for hyperlocal distribution. Furthermore, in case of a failure of the hyperlocal server, fall-back mechanisms have to be considered. RFC 7706 bis will also examine existing running code. One example for running code is the local host project by Wes Hardacker, which allows people to set up local resolvers with the root zone to "serve yourself".

More privacy steps: lifting QNAME minimisation from experimental to standard

Besides the hyperlocal concept, another existing technology driven forward via a bis-version of an RFC is QNAME minimisation (RFC 7816). Most DNS software now offers it (Bind is announcing it for the next weeks, Knot and Unbound offer it in their current versions as an option). With QNAME minimisation, no full queries will be sent up to the root zone but only those for TLD zones, thereby removing the possibility for root servers to keep track of individual queries.

There was broad agreement that making QNAME minimisation a standard rather than an experiment was a logical step. Again, Paul Hoffman (ICANN) is preparing the bis-document version. Hoffman pointed out that given that DPRIVE work (privacy-securing resolver to authoritative server work) was lagging behind heavily, pushing for QNAME minimisation made sense.

While from the outside, the various moving parts of the DNS evolution seem to be competitive, at least in some aspects, experts think that overlapping the various parts (DoH, DoT, QNAME minimisation, Hyperlocal) addresses different situations and could be complementary. The only problem in this regard could be the effort (and/or cost) of tweaking one's DNS to use all mechanisms, which might be rather daunting and would only be an option for larger organisations/companies.

Working groups

DNSOP Working Group

The biggest ongoing work in the DNSOP WG – not counting the DoH and DoT discussions in the background – is about how to ease the use of application/service specific addresses in a DNS record.

ANAME or http-minimal record

As it is considered that Service Resource Records (SRVs) are not easy enough to use from a web-perspective and that the existing CNAME records are not flexible enough (because they do not allow for additional addresses to be placed at the apex) the WG is looking for a solution. Two proposals were discussed in Bangkok. One is to create a new resources record called Address specific DNS aliases (ANAME). In looks ANAME is similar to CNAME, thought at the same time it is designed to be at the APEX.

Ray Bellis, ISC, now proposes another resource record which he calls a dedicated "minimal HTTP" resource record type. It should "facilitate redirection from the domain name portion of an HTTP(s) URI to the server hostname and thence to A or AAAA records". Contrary to ANAME it will replace CNAME completely.

The discussion on these issues is still ongoing and has so far been inconclusive.

Blockchain marrying DNS - DNS DID

One forward-looking topic in the DNS WG (also presented in the Decentralized Internet Infrastructure Research Group) is about what the DNS could offer to blockchain providers.

With regard to interoperability, the W3C has already started work and provided a URI scheme that allows unified addressing without the need of a central registration. The Decentralized Identifiers (DID) provides a naming convention similar to Universally Unique Identifiers (UUIDs). The difference is that DIDs can be resolved like URLs or dereferenced to a standard resource describing the entity and, unlike a classical URL they typically contain cryptographic material, enabling authentication of the entity responsible for the resource (did:example:123456789abcdefghi). According to the W3C draft, each DID contains at least "cryptographic material, authentication suites, and service endpoints". An experimental DID-Registry that has been set up currently lists a dozen Blockchain providers, including BitCoin (did:btcr: did:stack:), Etherum (did:csnt, did:erc725, did:uport) and Sovrin (did:sov:).

According to Alexander Mayrhofer (nic.at) the contribution of the DNS can be to allow easy (and global) addressability, as URLs are no easier to memorize or to read than Blockchain hashes. With RFC 7553 the technology is also already in place, allowing "URI Resource Record types". An update to the RFC only needs to add DIDs as a new type (_did.example.net. IN URI 100 10 "did:sov:1234abcd"). Linking DIDs to email addresses is also possible if a client asks for a DID instead of an OPENPGPKEY-record (see section 5 RFC 7929). With regards to potential loss of privacy/anonymity, Mayrhofer said that DID could be made for some blockchain applications that need to be found and public. He also noted that offering the technology as open standard based would prevent the potential creation of proprietary solutions.

Running code with a resolver is available here.

DNS ongoing work

Other ongoing work at the DNSOP is to extend the use of the Time to Live of a DNS resource record "in the exceptional circumstance that a recursive resolver is unable to refresh the information." By using stale data, outages of a server can be bridged. The proposal from Warren Kumari (Google) was welcomed and is already a WG document.

KSK Roll Talks

There will now be talks with dozens of communities (it was not up to the DNS community alone) to decide on the way forward for future rollovers. Several options could be discussed, including rolling at regular intervals – the ICANN CTO recently called a three-year interval a sensible time frame in a CENTR blog interview. At the same time there were also thoughts that a long-term key with one or several standby keys for emergency rolls could be an option.

Decent statistics were still lacking. Discussions during the meeting included:

- Whether there should be a back-up standby key (in an emergency, prepublication would not make sense anyway), Wes Hardacker warned that especially for distributed software, it might be beneficial to include prepublication and to have several keys in stock
- what changes should be made to RFC 5011 (automatic rollover)
- what changes should be made to RFC 8145 (Geoff Huston's question) to allow for better metrics (Signaling Trust Anchor Knowledge in DNS Security Extensions)
- when should algorithm rolls be considered (would it need more than one additional KSK rollover before an algorithm rollover should be considered?)
- how well did the outreach work for the rollover?

Geoff Huston warned that the KSK rollover was not as painless as portrayed by ICANN. He found problems in 75 networks (whose size ranged from 25 to half a million hosts), and potentially up to four million users experiencing issues in total. Huston confirmed that the biggest issue was experienced by the Irish network provider EIR.com. According to his figures, in all but three cases, networks fixed the problems themselves. Three networks just stopped validating. Huston warned that in older resolvers, systems could not be caught by the telemetry used. Hoffman reiterated that practically nobody had come to ICANN to complain or present problems they experienced during the rollover.

TLS WG: The end of the DNSSEC chain extension?

After a thunderous debate, in its second session in Bangkok the TLS WG moved to stop work on the DNSSEC chain extension proposal, after all the representatives of browser companies declared they would not implement this. As a result of the presentation given by TLS Co-Chair Sean Turner on the continuous back and forth of the work on the draft proposal, the WG took the draft out of the list of active WG documents. Turner acknowledged that the failure to complete the document was "not our finest hour". During the debate Wes Hardacker, in charge of the mentoring program for new members to the IETF said that 90% of complaints over toxic debates pointed to the TLS WG. The "toxic nature" of the debate was said by some as making them stop participating in the related discussions of the WG.

Last Call and Back

Since IETF93 the TLS WG has worked on a "new TLS extension for transport of a DNS record set serialized with the DNSSEC signatures needed to authenticate that record set" (DNSSEC chain extension). The authors are Melinda Shore (Fastly), Richard Barnes (Mozilla), Simon Huque (SalesForge) and Willem Toorop (Nl.net Labs). The idea of the proposal, which had a 2012 predecessor by Adam Langley (Google), was to allow TLS clients to perform DANE authentication of a TLS server without performing additional DNS lookups, thereby avoiding latency and last mile issues of DNSSEC.

Having already nearly passed the "IETF last call" earlier this year, the concern over a possible downgrade attack was raised when the IESG started to send back comments. The concern according to Turner is that "absent whitelists, a client misdirected to a server that has fraudulently acquired a public CA-issued certificate for the real server's name, could be induced to establish a PKIX verified connection to the rogue server that precluded DANE authentication".

Chairs have tried to push the document over the last hurdle several times, with the most recent debate trying to figure out if pinning could lift the concerns of the attack, and if changes in the focus would help finalise it. However, at this point it does not seem like it will be possible to reach a consensus in the WG anymore. During the Bangkok debate, some of the participants made harsh comments toward those still trying to fix problems (including Victor Dukhovni, OpenSSL Foundation and author of a DANE operational RFC, and Nicolas Willians, consultant at Cryptoconect). David Schinazi (Apple) not only said that Apple would not implement this, but also asked to "please kill this because it's wasting the working group's time". Schinazi later tried to take a step back, underlining that there still remained the informational RFCs.

After the WG removed the document from the TLS WG, it remains to be seen how the proponents will react and if they will try to come back with a changed draft or rather publish the document as an individual document.

QUIC Spin Bit finally accepted

After another hour of discussions, the QUIC WG finally decided they would include the Spin Bit in the QUIC Standard, version 1.0. The Spin Bit allows network operators (and others) to measure latency and to troubleshoot, according to its advocates. With the Spit Bit set, the server and client flip the bit when it reaches them, allowing the server to measure run times.

For many months, the WG had fought back against adding this extra Bit, which was asked for by network operators and some intelligence agencies (National Cyber Security Center), because QUIC encrypts additional parts of the transport headers, thereby taking out meta data to be used for the monitoring/ surveillance of traffic.

Several changes were made in the Spin Bit proposal to alleviate concerns over surveillability. First, each client and server will have the option to decide whether to have it set or not. In order to prevent those who choose not to set it from sticking out (thereby attracting attention), operators have to ensure that the Spin Bit is not set on all their connections. They must have an "anonymity set".

On of the problems with these measures is that it is unclear if operators will implement the anonymity set or not. Representatives from browser companies (Google, Mozilla) and platforms (Facebook), plus providers Fastly and Protocol Labs announced that they would not implement the Spin Bit, at least for now. Microsoft, Apple and Broadcom announced that they would deploy the Spin Bit. The Spin Bit remains a trade-off according to experts: with the trouble-shooting and latency measurements, attacks could be prevented or countered, making users more secure. At the same time the extra bit gives a tiny bit more information about an endpoint than necessary.

The QUIC WG lags behind its original, highly ambitious time plan, but the final QUIC version one RFC should be finalised early next year.

SUIT: One or more formats?

The SUIT ("software updates for the Internet of things") working group is still finalising its architecture and information model documents, and is slightly lagging behind the milestones set out. In Bangkok the most important topic addressed was the question on whether the WG should adopt one, and only one, data format for the manifest, or whether it should allow several, as long as they use the same data model.[1]

1 Manifests are "a bundle of meta data about the firmware for an IoT device, where to find the firmware, the devices to which it applies, and cryptographic information protecting the manifest.

The differences between the two formats presented at IETF103, one by ARM developers Brendon Moran and Hannes Tschofenig, the other by Martin Pagel from Microsoft, lays mainly in the protocol used to express the format. Moran and Tschofenig propose to use the Concise Binary Object Representation (CBOR) which claims to "already be optimized to be small code size, small message size and extensibility without the need for version negotiation." Pagel proposes a simple binary txt format instead, and gives a glimpse of the differences between these two candidates.

Pros and Cons vs CBOR based Format

"CBOR makes it easier to handle and/or skip optional or new fields, whereas a binary structure requires a versioned structure to introduce new fields, which adds complexity to the implementation. However, the binary structure has the advantage that it can be loaded into memory directly without the use of a parser and therefore the installer code is much simpler or smaller. As installers are a common source of bugs and vulnerabilities, simple code is usually considered more secure. It addresses Section 3.6/7 of the architecture document (Small bootloader and parser) quite well. Also, the separation of image URIs allows for a much smaller manifest and therefore reduces memory requirements. A basic device may not be able to support many options anyways and such devices are more space constrained; the binary format may be a better fit. A more sophisticated device may offer more options and may use CBOR for other purposes anyways, then the currently proposed format may be more suitable."

A preliminary hum taken by the WG showed that a majority favours keeping the number of data formats at one, which is rather an ambitious goal. At the same time, nearly as many indicated that they "need more information". The discussions will therefore continue. A third proposal from the Fraunhofer SIT was neither presented not discussed in Bangkok.

One question raised by Gurshabad Grover from the HRPC RG was if it was enough to secure the path: the discussion on the additional securing of data was inconclusive.

Other issues discussed in the WG was the liaison on IoT with the ITU. The WG agreed to keep terminology as aligned as possible between SUIT and SG17, which is currently finalising its own document on IoT.

Meanwhile and related to the work on IoT, CIRA is pursuing the standardisation of its secure gateway project, but still seems to be looking for the right working group/place in the IETF.

Research Groups

Is the SMART RG a smart idea?

After having lost the fight for a static key for the new TLS 1.3 (see the IETF 101 Report), law enforcement and a number of companies have tried to re-gain lost ground. One attempt to do that is the initiative by the British National Cyber Security Center (NCSC) to establish a group dedicated to research on how threats might be detected in encrypted traffic. The Stopping Malware and Researching Threats Research Group (or SMART RG) will investigate how cyber-attack defence requirements can be met in a world of encrypted data," the draft charter for the SMART RG reads.

During a side meeting (a new session format at the IETF 103) chaired by Kirsty Paine (NCSC) and former Security Area Director Kathleen Moriarty (Dell). Paine underlined that the working group did not intend in any way to weaken new security standards developed at the IETF, including encryption. Nevertheless, the research group intended to "research the effects, both positive and negative, of existing, proposed and newly published protocols and Internet standards on attack defence."

Attacks could be malware, phishing, DDoS and also, as Moriarty said, pervasive monitoring. During the session Paine listed the endpoint detection capabilities and limitations, threat detection in encrypted traffic and metrics for goodness and badness as some of the research topics.

Research group members, including law enforcement or intelligence agencies, can bring case information to illustrate issues, and the RG can comment on standards under development, offer alternatives to reach better attack defence levels or even propose solutions. According to the draft Charter, "within the first year, the research group aims to:

Survey existing attack detection methods and determine the relative effectiveness of these methods against different attack defence threats (e.g. phishing, DDoS, spambots, C&C, endpoint malware);

Publish case studies of historical attacks and make recommendations where attacks could have been stopped more quickly, or even prevented.

Publish an Informational RFC, titled: "Important Attack Defence Considerations for Protocol Design and Deployment".

Attack Defence Considerations?

The RFC "Important Attack Defence Considerations for Protocol Design and Deployment" pretty much mirrors RFC 8280 (and the related guideline document under consideration), and there seems to be an expectation that the law enforcement/defence side will then also be able to give their advice.

Paine also underlined that they did not intend to overlap with the work the Security area does in advising WGs on potential risks and security issues. However, one might wonder if the group decided they needed to step up their game in terms of law enforcement (public security?) considerations, given the previous security, privacy and human rights efforts.

Reactions from participants during the well-attended side meeting (25-30 people) were mixed. Former Security AD Stephen Farrell recommended that the Charter should not to be too ambitious, but that at the same time concrete results should be delivered quickly.

Bret Jordan, from the Office of the CTO at Symantec, applauded the initiative and said that the group would fill a large gap at the IETF and that "if we get the marketing right, there will be a lot of people".

The first official RG meeting is expected to take place in Prague (IETF 104). The group is also sending out a request for papers for the second CARIS (Coordinating Attack Response at Internet Scale) workshop in the next two weeks.

Standardisation forum shopping/attack on TLS?

Since the end of IETF 103 the mailing list of the group has been silent, with the exception of a pointer by Daniel Kahn Gillmore (ACLU) on research that looks into the risks of user tracking via TLS resumption (avoiding round trips by resuming TLS sessions with already visited hosts).

Controversy can nevertheless be expected to peak over an attempt at "forum shopping", with several proponents having published a version using static keys to break the end-to-end concept and allow for interception. The IETF meanwhile has asked ETSI, the standards body in question, to desist from calling eTLS TLS at all. The European Standardisation body ETSI just announced its TLS 1.3 variant (enterprise TLS, formerly multi-context TLS) for enterprise TLS (eTLS) which will allow data centre managers to keep TLS keys for the edge points under their control (at least). The rationale for the ETSI proposal is mainly to break the new TLS 1.3:

"Requirements - such as legal mandates and service agreements - exist for enterprise network and data centre operators and service providers, organisations, and small businesses to be able to observe and audit the content and metadata of encrypted sessions transported across their infrastructures [i.2]. The original TLS protocol standard adopted in 1994 and its subsequent versions up to and including TLS 1.2, provided for these capabilities [i.3] and [1]. The latest version of the protocol, TLS 1.3, does not provide for these capabilities [2]. Where these capabilities do not exist, this new encryption protocol could be blocked altogether at the enterprise gateway, forcing users to revert to older, less secure protocols. The present document is one of a series of implementation profiles that, to achieve these required capabilities, puts the enterprise operators and users in control of the

EXAMPLE 2: Middlebox B decrypts the traffic in real-time to provide application health monitoring, but also stores the encrypted packets so they can be decrypted at a later date for compliance and auditing purposes.

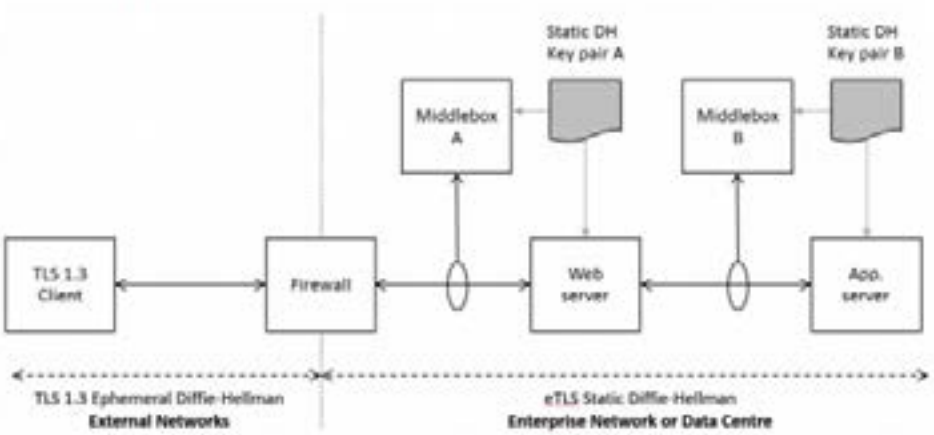

Figure 4.1: eTLS architecture with enterprise servers

4.2.2 eTLS with enterprise clients

Figure 4.2 depicts the eTLS implementation architecture when used with enterprise clients. TLS connections to servers that are external to an enterprise network may be made using TLS 1.3 [2], using forward secrecy and enhanced protections.

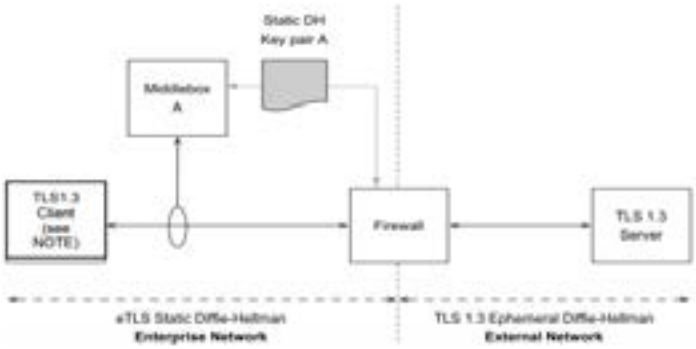

Figure 4.2: eTLS architecture with enterprise clients

access to their data for cyber defence and prevents unauthorized access. It sets forth a 'Profile for an enterprise network and data centre access control' called eTLS that meets several desired capabilities for the Middlebox Security Protocol MSP [i.1]."

In essence, with eTLS in place, TLS 1.3 is terminated at the Firewall of the network, and either the Firewall or an internal web server act as a TLS 1.3 proxy between the network client and an external application or an external client and the application inside the network.

The proposal essentially takes on what was discussed but failed to receive consensus during the IETF TLS WG discussions. At the same time, it limits visibility. clients who have not implemented eTLS visibility will not know that middleboxes that decrypt their traffic are involved, and that naturally, they will also not "receive, if requested, validation of identity by each middlebox."

What remains an open question is how the IETF will react to the forum shopping on TLS, for which it claims IP rights and change control. While the ETSI Cybersecurity Technical Committee (TC Cyber) promised "not to use the name TLS apart from referring to the IETF standards" in a letter, the published document clearly calls the new ETSI standard an "implementation variant" of TLS. ETSI also points to similar standards allowing middlebox decryption at the ITU.

Besides this, ETSI challenges the fact that the IETF can claim copyright on TLS and refers to related technology which, according to ETSI predates IETF's TLS standard suite. The IETF Security Area rejected this statement in their liaison statement on 5 December.

IETF has lost earlier fights on T-MPLS, and ETSI has previously been used by EU law enforcement in standards shopping events, especially with regards to Lawful Interception (LI) of communication networks.

Quantum Research Group

The IRTF is about to charter a Quantum Research Group which, like the new SMART RG will officially meet for the first time at IETF 104 in Prague. At a very well-attended (100 people) preparatory session, RG initiator and future Co-Chair, Rod van Meter (Keio University) explained the rationale for the research group. While work toward the Quantum Internet is well underway, the physicists working on entanglement and key exchange are lacking knowledge in network engineering. Furthermore, the Quantum researchers acknowledge that with regards to applications they need to advance in deciding "what we would *do* with a Quantum Internet" and how to develop a multi-party system (instead of simple point-to-point transfer systems which are currently under development). A vision for a roadmap for Quantum networks has been described by the other Co-Chair, Stephanie Wehner (TU Delft) and two colleagues here. The RG hopes to become a focal point for quantum networking standardisation and also intends to consult the IETF on Quantum crypto.

According to the proposed charter, the research group will work on:

- **routing**: there have been a number of proposals, including a couple in the last six months or so, and so there will need to be an assessment of which routing schemes are appropriate under which circumstances
- **resource allocation**: some of the routing proposals seem to include a notion of the dynamic traffic on the network, but this distinction needs to be defined clearly
- **connection establishment:** what does a request look like (semantics more than syntax) as it propagates across the network?
- **interoperability:** given than different networks are currently being designed and built, how do we ensure the development of a long-lived internetwork?
- **security**: are quantum repeater networks inherently more or less vulnerable in operations than classical networks?
- **design of an API** that will serve the role that sockets play in classical networks

Certificate Transparency

In two very interesting presentations (during the IRTF open meeting and MAPRG) researchers from the ICSI California (Johanna Aman) and the University of Hamburg (Matthias Wählisch) presented statistics on the evolution of TLS and certificate transparency (CT), illustrating the implications of exposing certificate DNS names from the perspective of security and

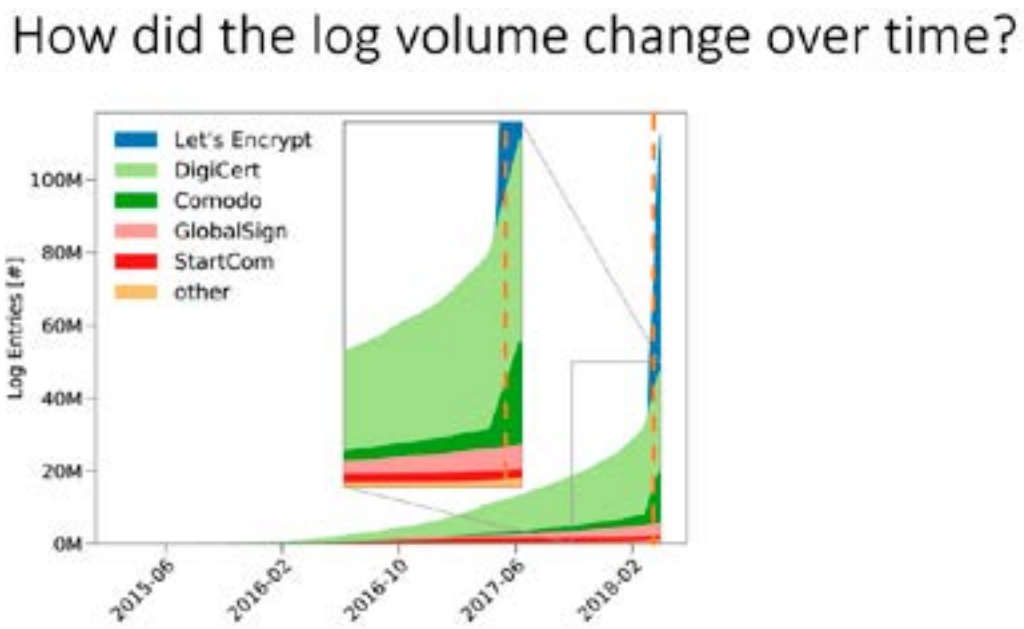

privacy. Though they found that an exponential growth of certificates had been listed in transparency logs, and that website support for CT represents 33% of established connections, they also found that it took only one hour before there were the first DNS lookups for domains they had set up with the certificates listed. In most cases the researchers could not find who had scanned them (there was no information on rDNS, WHOIS or the website). One scanner requested fast A/AAAA records and scanned 30 ports. While certificate transparency helps to find phishers, information leakage remains a problem. The researchers also found that only a few logs held all log entrances.

Researchers also commented on added privacy via encryption. They see it hurting their measurements but think that on balance the added encryption is the right way to go.

HRPC struggling over own procedures

Beside the continued talks on how the HRPC should advance its work, Co-Chair Ari Doria warned that the group must not act as if it were a full-fledged "Directorate" of the IETF, as it was not – the group had invited Arthit Suriyawongkul from the Thai Netizen Network to talk about the freedom of assembly and technology. The Netizen Network monitors current legislative action in their country, namely data protection legislation (which allows for a lot of exemptions on national security, police, insurance companies and others) and cyber security legislation.

According to Suriyawongkul, an issue from the NGO's point of view is that often, the perceived security does not match the real security offered by protocols. In his presentation, he discussed the right to online association, linking it to both freedom of expression (basic, individual) and privacy (conditional to exercising other rights), and also pointed to new ways protestors can be attacked. In several countries police have started to play copyrighted music during protests, resulting in automatic take-downs of the streamed protest marches (and speeches) from YouTube for copyright reasons.

Trojan horse in HRPC?

A proposal that is still rather vague on how filtering should be perfected using IETF protocols, based on the concept that IETF should not cherry pick rights, but defend all human rights, was brought to the HRPC goup by Nalini Elkins, Enterprise Data Center Operators. Although Elkins focussed on various filtering examples and rationales in the draft, during her talk she tried to drive home the rationale that filtering was necessary to avoid human casualties. She argued that the human right to life is at odds with the right to freedom of expression and to privacy. Elkins belongs to those calling for a static key in TLS 1.3. During the session, RG Chair Avri Doria pointed out that the "filtering draft" should focus on the description of the filtering landscape and not become any kind of operational draft.

IETF News

Participation in Bangkok was very low, with only around 850 participants. IETF Chair Alissa Cooper, who was presenting the stats, argued it was a natural up and down.

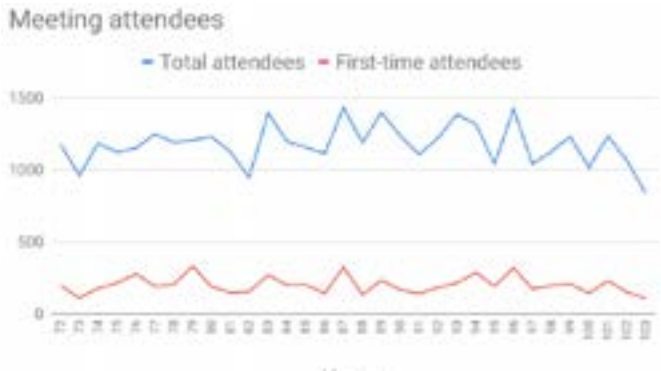

From the various experiments to attract new people, the hackathon seems to be the most promising, with Cooper reporting that there were quite a number of people travelling to the IETF only to attend the Hackathon.

Without much fanfare the IETF has set up its LLC organisation, which was legally created on 27 August 2018, and it has assumed the IAOC's responsibilities. The IETF Executive Director replaces the IAD position, and the position will be filled by the NomCom, with Portia Wenze-Danley currently acting as interim Executive Director. The IETF LLC Board of 5 Directors (1 chosen by IESG, 1 chosen by the ISOC Board of Trustees, 3 chosen by NomCom) will assume the role of the IAOC. The full Board will be announced at IETF 104. The current interim Boad is made up of Glenn Deen, IAOC Chair (LLC Chair); Alissa Cooper, IESG Chair; Ted Hardie, IAB Chair and Gonzalo Camarillo, ISOC BoT Chair.

IAB and technical Plenaries

The open Friday for non-WG side-meetings experiment did not completely work out, as many side-meeting organisers had chosen to have the side-meetings before Friday. There was also a complaint by ISOC Chair Andrew Sullivan, that the technical plenaries nearly had to be abandoned. Nonetheless, as Sullivan said, the technical plenaries allowed the IETF community to have cross-WG talks about interesting and current technical developments. The program committee for the tech plenaries is looking for additional people, but will have technical plenaries again in the future, according to IAB Chair Ted Hardie.

The IAB has published several reports on workshops that took place years ago (RFC 8477, Report from the Internet of Things (IoT) Semantic Interoperability (IOTSI) Workshop and RFC 8462, Report from the IAB Workshop on Managing Radio Networks in an Encrypted World (MaRNEW)). There is now a discussion on how to deal with reporting back. With regards to transparency, the IAB reacted to complaints and has opened up its meetings to observers. Furthermore, the agendas for IAB teleconferences will now be published. An interesting read is the letter the IAB sent to the Australian legislator in one of its rare political statements with regard to Australia's proposed Assistance and Access Bill (that will break encryption).

New appointments

Tim Wicinski - Community Coordination Group (advising the IETF Trust)

Ole Jacobsen – reappointed by ICANN NomCom

Sarah Banks, Tony Hansen, Adam Roach, Peter Sant-Andre, Robert Sparks, Christian Huitema – members of the RFC Series Oversight Committee

Calls out (soon)

The IAB is searching for a new IRTF Chair as Alison Mankin is stepping down next year.

Volunteers for ICANN Technical Laison Group.

Board of Trustees ISOC

CENTR is the association of European country code top-level domain (ccTLD) registries, such as .de for Germany or .si for Slovenia. CENTR currently counts 55 full and 9 associate members – together, they are responsible for over 80% of all registered domain names worldwide. The objectives of CENTR are to promote and participate in the development of high standards and best practices among ccTLD registries. Full membership is open to organisations, corporate bodies or individuals that operate a country code top level domain registry.

Rate this CENTR Report on IETF103

(Thank you for your feedback!)

CENTR vzw/asbl
Belliardstraat 20 (6th floor)
1040 Brussels, Belgium
Tel: +32 2 627 5550
Fax: +32 2 627 5559
secretariat@centr.org
www.centr.org

To keep up-to-date with CENTR activities and reports, follow us on Twitter, Facebook or LinkedIn